


CONNEXIONS



The Interoperability Report

May 1991

Volume 5, No. 5

*ConneXions —
The Interoperability Report
tracks current and emerging
standards and technologies
within the computer and
communications industry.*

In this issue:

Components of OSI: CO/CL Interworking.....	2
NIST OSI Routing Testbed.....	7
The X Window System and OSI VT.....	8
Resource Discovery.....	12
Book Reviews.....	21

From the Editor

Our series *Components of OSI* continues this month with a look at the problem of connection-oriented (CO) and connectionless (CL) transport interworking. In Europe, transport service is typically realized with a low overhead version of OSI Transport known as Class 0 or TP0 running on top of X.25 which offers end-to-end data integrity. In the US, connectionless networks are much more common, and the corresponding transport service used is Class 4 or TP4. This transport class has reliability built into it (much like TCP) and thus it "does not care" about the underlying network service.

You could argue forever about which of these two styles of networking represents the "better" solution, but we'll skip the religious debate and get right to the problem: How do you connect a network based on the TP0/X.25 model to one running TP4/CLNP and make their respective end-systems interoperate? This problem must be solved if OSI networks are to become more than "regional solutions." We asked Marshall Rose to describe the problem and give an overview of the efforts which are underway to solve it. Marshall has been a frequent contributor to *ConneXions* ever since its first publication back in May of 1987. We are very grateful to him and the countless other authors who have helped us develop a total of fifty editions since that time.

The OSI virtual terminal protocol was described in our January 1991 issue. Since The X Window System offers some features which are similar to that of VT, we thought it would be interesting to find out what exactly the relationship is between these two protocols, and asked Graham Dixon of Churchill College, Cambridge to explain. His article appears on page 8.

Michael Schwartz describes means by which users can discover the existence of *resources* (such as network services, documents, and people) in a large internet environment. This problem is important because it is an enabling technology for the larger issue of supporting distributed collaboration among many interrelated individuals across administrative boundaries.

Also in this issue, you'll find a couple of book reviews, and some news about an OSI Routing Testbed that is being supported by the National Institute of Standards and Technology (NIST).

We'd like to remind you once again about our 1991 *Internetworking Tutorials Program*. Your last chance to take these courses before INTEROP 91 will be May 13-16 in Dallas, Texas. For more information, or to register, call 1-800-INTEROP or 415-941-3399.

ConneXions is published monthly by Interop, Inc., 480 San Antonio Road, Suite 100, Mountain View, CA 94040, USA. 415-941-3399. Fax: 415-949-1779. Toll-free: 1-800-INTEROP.

Copyright © 1991 by Interop, Inc.
Quotation with attribution encouraged.

ConneXions—The Interoperability Report and the *ConneXions* logo are registered trademarks of Interop, Inc.

ISSN 0894-5926

Components of OSI: CO/CL Interworking

by Marshall T. Rose Performance Systems International, Inc.

Introduction

The OSI transport service [1] may be realized through a variety of transport/network protocol combinations. Regrettably, few of the combinations actually interoperate with each other. For example, an end-system running TP4 and CLNP cannot interwork with another end-system running TP0 and X.25. Hence, even if all OSI-capable end-systems enjoyed full-connectivity, they would not be able to uniformly interoperate.

In July of 1990, at the request of the Federal Networking Council (FNC) and the International Collaboration Board (ICB), a workshop was convened to examine solutions to this problem. It was hoped that this would allow the Internet and RARE communities to deploy different "styles" of OSI and still achieve interworking across enterprises. The scope of the workshop was to address interworking problems encountered when two end-systems attempt to interoperate using a common application protocol, e.g., the *OSI Directory* (X.500). It was outside the scope of the workshop to worry about the interoperability problems between two end-systems using different application protocols, e.g., one end-system using the OSI file service (FTAM) and the other end-system using the Internet file transfer protocol (FTP). In this case, an *application gateway* technology should be used.

Application use of End-to-End Services

To appreciate the direction taken by the workshop, we must first understand how an OSI application makes use of OSI end-to-end services. Conceptually, when an OSI application wishes to establish an association, that application identifies an *application entity* that provides the service desired for communication. The application entity identification is given to the OSI Directory and the corresponding *presentation address* is retrieved. This presentation address consists of a *presentation selector*, *session selector*, and a *transport address*. When the association is to be initiated, there are two parameters of interest: the presentation address as provided by the Directory, and the communications *Quality-of-Service* (QoS) as desired by the application. The first identifies the location of the desired service, the second identifies the characteristics of the association to be established with that service.

After passing through the highest layers, the transport address, consisting of a *transport selector* and one or more network addresses, is given to the transport service, and a request is made to establish a transport connection. The local transport entity now follows three steps in order to establish the connection:

1. The entity looks at each network address and decides which mode of network service, *connection-oriented* (CONS) [2] or *connectionless-mode* (CLNS) [3], will be used for the address. At present, there is no standard method nor a set of agreements for making this determination; in some implementations, the determination is made on the basis of NSAP prefixes, with this information being configured by the system administrator.

Based on the derived network service and the desired QoS, the local transport entity selects a transport protocol. That is, for each network address in the transport address, the entity selects a combination of a transport protocol and network service, referred to as a *TS-stack*, that will be used to establish a transport connection to that address.

2. The network addresses are then ordered, based on the desired QoS and the "closeness" of the network address. Again, this decision is a local matter.

Suppose, for example, a transport address contained two network addresses, each implying use of a CONS. One of the network addresses might reside in a private network, whilst the other address resides in a public data network. For economic reasons, the local transport entity might prefer to try the private network first.

3. For each network address, the local transport entity starts the protocol machine for the TS-stack associated with that address. This results in a transport protocol and underlying network service being invoked. Once a transport connection is established, the remaining network addresses are ignored.

Of course a TS-stack can be realized using non-OSI protocols. For example, the RFC 1006 method defines a *transport service convergence protocol* which smoothes over the differences in the services offered by the OSI transport service and TCP [4]. This is known as the RFC 1006/TCP TS-stack. If such a protocol is properly defined, then the OSI upper-layer infrastructure (session, presentation, and application layers) running above is unable to tell that it is not running in a "pure" OSI environment.

Problems in realizing The Model

Looking back at the first step, the entity must establish a binding between each network address and a TS-stack.

Suppose that the end-system on which the transport entity resides offers only a subset of the TS-stacks implied by the transport address. For example, suppose that there are four network addresses, two requiring use of the CONS, and the other two requiring use of the CLNS. If the initiating end-system supports only the CONS, then any network address which requires use of the CLNS cannot be reached from that end-system. That is, the local transport entity must intersect TS-stacks derived for the network addresses with the TS-stacks supported by the local end-system. Thus, in this first step, only a subset of the network addresses may be suitable for use on the initiating end-system.

The problem, of course, is that the intersecting subset may be empty! From a purist's perspective, interworking can not occur, and the local transport entity will immediately generate a transport disconnect.

In exploring this problem, it is natural to ask how often this situation arises. The answer is simple: in a homogeneous OSI environment, say a CL-mode LAN (e.g., an 8802 network) or a CO-mode WAN (e.g., an X.25 network), this problem should never arise. However, whenever different OSI environments are interconnected, this problem usually results.

Consider the simplest example: a site has an 8802-based subnetwork running CLNP, TP4, and OSI applications. All of the end-systems in that environment implement the TP4/CLNS TS-stack. Some time later, one of the end-systems on that subnetwork is attached to an X.25-based subnetwork. For brevity, term this end-system "dual." On the dual end-system, another TS-stack is added, e.g., TP0/CONS. The other end-systems are not modified since they continue to have a single point of attachment, which supports only the CLNS. Now, observe that within the original 8802-based subnetwork, all end-systems, including dual, continue to interoperate with one another.

CO/CL Interworking (*continued*)

Also observe that the dual end-system can interoperate with any other end-system directly attached to the X.25-based subnetwork. However, note that it is unlikely that any of the other end-systems in the 8802-based subnetwork can interoperate with an arbitrarily chosen end-system in the X.25-based subnetwork.

In order to appreciate the basis for the approach which follows, it is necessary to introduce one additional term, the OSI *community*. An OSI community is a collection of end-systems connected together and sharing a common TS-stack. More technically, a community is defined in terms of connectivity and a TS-stack.

So, given two OSI communities which have an intermediate-system in common, but have different TS-stacks, can arbitrary end-systems in those two communities interoperate? First, note that the CONS and the CLNS do not interwork. Hence, if the two communities support only different modes of network service, then they cannot interoperate. Second, note that even if two communities share a network mode in common, then all intermediate-systems must also support that same network mode. For situations in which direct interworking is not possible, a transport-layer relaying approach has been suggested. Because they exist outside the scope of OSI, the theory and practice of transport-layer relays are poorly-understood.

Transport-Service bridges

The motivation behind transport-layer relaying is to observe that all TS-stacks share one thing in common: they all offer the OSI transport service. Thus, a new entity is introduced, residing above the transport service, termed a *transport service bridge* (or *TS-bridge*), or *CO/CL-gateway*. The TS-bridge is purposefully naive as to TS-stacks or the transport protocols and network services which compose them. Rather, the TS-bridge knows only how to invoke the OSI transport service, which is offered by all TS-stacks, regardless of their composition. In pictorial form:




Figure 1: A Transport-Service Bridge

The function of the TS-bridge is simple: upon receiving a transport connection indication, the TS-bridge initiates a second transport connection, to the actual destination address. If the second connection is established, then the TS-bridge accepts the first transport connection. From this point on, any data received on one transport connection is simply sent on the other transport connection. When a disconnect occurs on one of the transport connections, the TS-bridge disconnects the other transport connection. Of course, if the second connection is not established, then the TS-bridge will simply disconnect the first transport connection.

It should be noted that transport-layer relays suffer from (at least) four weaknesses:

1. A transport-layer relay maintains state for its two existing connections, and is therefore a single point of failure. For example, if the relay fails, the transport connections between the end-systems will fail, even though both end-systems are operational and an alternative path is available.
2. Use of a transport-layer relay defeats the end-to-end integrity property of the TS-stack. Note that user-data passes through the relay in transit between the two TS-stacks. This data might be corrupted if the relay is faulty.

Similarly, use of a transport-layer relay defeats any transport-level encryption mechanisms as the data appears in the clear inside the relay. (Of course, encryption could occur at a higher layer to retain privacy.)

3. Use of a transport-layer relay may introduce additional variability in round-trip times due to buffering in the relay. (The implications of this effect are not known.)
4. Finally, depending on how use of the transport-layer relay is integrated with the end-systems, end-to-end addresses may not be carried transparently. For example, in the short-term, the responding end-system sees the network address of the transport-layer relay as the calling address, instead of the address of the actual originator.

Towards a solution

The explanation above did not indicate how a transport connection was redirected to a particular TS-bridge. Depending on the approach taken, varying levels of transparency can be achieved in the end-systems.

One solution requires knowledge of the TS-bridge in the initiating end-system; further, such a solution requires that the initiating end-system act in a non-standard fashion in order to establish a connection when using a TS-bridge.

In contrast, another solution might rely on a rich OSI network-layer infrastructure so as to achieve the "ES-transparency" effect: no local knowledge of TS-bridges should exist at an end-system; further, use of a TS-bridge should not result in non-standard behavior at an end-system.

In the best solution, there is a single mode of OSI network service which is truly ubiquitous. In this case, a single community exists and interworking is achieved through the use of network-layer relays. In preparation for this long-term scenario, technology must be identified and perhaps incrementally advanced to promote a homogeneous network service. In the meantime, a large TCP/IP-based community exists and a TP0/CONS community is growing. Some interworking requirements exist today and these requirements are expected to increase.

This suggests a short-term solution to address immediate requirements, an intermediate-term solution applicable as the TP0/CONS community grows large, and a long-term solution applicable once two large OSI communities, one CO-mode and the other CL-mode, exist and have interworking requirements.

continued on next page

CO/CL Interworking (*continued*)

Thus, an approach towards the solution consists of two parts, and two companion memos have been written:

- In the short-term, one must rely on TS-bridges to provide connectivity between non-internetworking communities. The first companion memo, [5], describes the operation of TS-bridges in such an environment.
- However, even in the long-term, situations will arise in which both network services are required. In this case, TS-bridges are still necessary. The third companion memo, [6], describes the operation of TS-bridges in such an environment.

Acknowledgements

The CO/CL Interworking Workshop met on July 24–26, 1990, and was co-chaired by Les Clyne of the Joint Network Team and Phill Gross of the Corporation for National Research Initiatives, and had participants from several European and North American concerns. The workshop produced four documents, based on the contributions presented. At a subsequent tele-meeting, these documents were revised in the light of recent standardization activity, and one document proved unnecessary. The remaining three documents which have been made available as Internet-Drafts, will soon be published as RFCs. This article is based on the overview document produced by the workshop [7].

References

- [1] International Organization for Standardization, "Information processing systems: Open Systems Interconnection, Transport Service Definition," International Standard 8072, June 1986.
- [2] International Organization for Standardization, "Information processing systems: Data Communications, Network Service Definition," International Standard 8348, April 1987.
- [3] International Organization for Standardization, "Information processing systems: Data Communications, Network Service Definition—Addendum 1: Connectionless-mode Transmission," International Standard 8348/AD 1, April 1987.
- [4] M.T. Rose and D.E. Cass, "ISO Transport Services on top of the TCP," RFC 1006, May 1987.
- [5] M.T. Rose (editor), "An Approach to CO/CL Interworking—Part II: The Short-Term—Conventions for Transport-Service Bridges in the absence of Internetworking," CO/CL Interworking Workshop, July 1990.
- [6] C. Huitema (editor), "An Approach to CO/CL Interworking—Part IV: The Long-Term—Conventions for Network-Layer Relays and Transport-Service Bridges in the presence of Internetworking," CO/CL Interworking Workshop, July 1990.
- [7] M.T. Rose (editor), "An Approach to CO/CL Interworking—Part I: Introduction," CO/CL Interworking Workshop, July 1990.

MARSHALL T. ROSE is Principal Scientist at Performance Systems International, Inc., where he works on OSI protocols and network management. He is the principal implementor of the *ISO Development Environment* (ISODE), an openly available implementation of the upper layers of the OSI protocol suite. He is the author of *The Open Book: A Practical Perspective on OSI* and *The Simple Book: An Introduction to Management of TCP/IP-based internets*, both professional texts published by Prentice-Hall. Rose received the Ph.D. degree in Information and Computer Science from the University of California, Irvine, in 1984. His subscriptions to *The Atlantic* and *Rolling Stone Magazine* are in good standing.

NIST OSI Routing Testbed

The current set of mature OSI routing protocols provide for significant dynamic, adaptive routing functionality. Early implementations of these protocols are beginning to emerge. The importance of the service provided by these protocols and the complex, multi-peer nature of their operation will necessitate means of testing conformance, demonstrations of multi-vendor interoperability, and methods for product evaluation before they can be confidently employed in large-scale networks and mandated for use in government procurements.

Objectives

To address these concerns, the National Institute of Standards and Technology (NIST) is establishing a cooperative laboratory for OSI routing technology. This effort has multiple objectives:

- Establish a cooperative research program with participants from industry, academia, and government interested in fostering conformance, interoperable OSI routing products.
- Provide an open testbed facility for OSI routing products.
- Foster mature commercially available OSI routing products.
- Research and develop methodologies and prototype tools to support conformance testing, interoperability testing, and product level evaluation of OSI routing technology.

Benefits

NIST believes that this program of work provides benefits to both the vendor and user communities. Experimental conformance and interoperability testing in a multi-vendor environment provides implementors valuable feedback, thus expediting the availability of product-level OSI routers. The research and development of testing methodologies will help fulfill NIST's responsibilities in the GOSIP program of insuring that means of assessing conformance to standards and multi-vendor interoperability are available. The development of product evaluation guidelines will help fulfill NIST's responsibilities in assisting government agencies in the evaluation, acquisition and use of emerging technology.

It should be noted that this effort will not result in any products being officially tested or rated. NIST's mission is to develop the means of testing. Of course, in this process we will perform experimental interoperability testing and conformance testing on implementations that participants have contributed to this effort. The results of this experimentation will provide feedback to NIST (in the development of testing methodologies), to standards communities (refinement of base standards, implementor's agreements, and user group profiles), and to implementors.

Project plan

NIST invites implementors, users, system integrators and other experts in routing protocols to participate in this effort. A project plan detailing NIST's program of work and ways in which interested parties may participate in this program is available for anonymous FTAM (ISODE, user: ftam, realstore=unix) and FTP from `osi3.ncsl.nist.gov` (129.6.48.108):

`./pub/doc/nist-routing-lab.ps`

To receive this document by e-mail (PostScript) or surface mail (hardcopy), or for further information contact:

Doug Montgomery

`dougmon@osi.ncsl.nist.gov`

NIST

Technology Building, B-217

Gaithersburg, MD 20899

Voice: +1-301-975-3630

Fax: +1-301-590-0932

The X Window SystemTM and ISO VT

by Graham Dixon, Churchill College, Cambridge, England

Introduction

There is currently under development a four-part *American National Standard* (ANS) for the X Window System Version 11. This work is being undertaken by ANSI Technical Committee X3H3, Computer Graphics, through its Task Group X3H3.6, Window Management. The full titles of the Standard and its parts are: *American National Standard for Information Systems—Computer Graphics—X Window SystemTM Data Stream Definition*;

Part I: Functional Specification

Part II: Data Stream Encoding

Part III: KEYSYM Encoding

Part IV: OSI Mapping.

On completion, it is the intention of ANSI to submit the American National Standard to ISO for ballot as an International Standard. Its status as an ANS will enable it to enter the ISO document cycle described in [1] at the level of a *Draft International Standard* (DIS).

The first three parts of the Standard are based on the X Window System Protocol Version 11 as specified for example in [2]. This article is not concerned with these parts, which give details of the X Window System itself. The surprise lies in the very existence of Part IV. This will enable the X Window System to be installed as an *Application Service Element* (ASE) in the OSI Application Layer, in a manner comparable with the ISO *Virtual Terminal* (VT) ASE described in [3]. Because of this linkage, functional standardization activities based on these X Window System standards will be undertaken in the three Regional Workshops NIST OIW, EWOS and AOW by the corresponding Special Interest Groups for ISO VT.

This article describes the usage of the X Window System as an ASE and considers its relationship with VT. The X Window System itself has been described elsewhere in this journal [4, 10], and an elementary knowledge of it is assumed.

The X Window System in OSI

The X Window System does not fit naturally into the structure of OSI. Its protocol combines the functionality of the OSI Session, Presentation and Application Layers and is designed to operate directly over a reliable byte stream such as is provided in OSI by the Transport Layer. However, it is not permitted by the OSI *Basic Reference Model* [5] to run an application directly over a transport-connection. The OSI mapping thus makes the minimum use of the OSI upper layers that is consistent with proper operation of OSI.

The Session and Presentation Services have been described in this journal in [6] and [7] respectively. Each of these services has a kernel functional unit that is always available, together with a number of optional functional units. To achieve minimum use of these layers, only the services of the kernel functional units are used. These provide to the Application Layer the facilities necessary for the establishment and release of a presentation-connection, together with the P-DATA presentation-service that provides normal duplex data transfer over such a connection. All other facilities of these layers, including expedited data transfer, resynchronization and the management of access tokens, are absent as they are provided by optional functional units. This level of service is thus superficially similar to that which would be provided by direct use of the Transport Layer. The conceptual level is however quite different, as will now be seen.

Titles and Addresses

The Application Layer of a real open system may contain a number of application-entities, each of which represents the OSI communication facilities of an associated application. A single application may be associated with more than one *application-entity*, each of which represents a different communication interface. The address provided in OSI for the establishment of a presentation-connection identifies a particular application-entity. At the client end of an X connection, such an address provides sufficient detail to identify the X Window System interface for an application that may have more than one means of access. At the server end, both the host and the particular X server are identified.

An application-entity is composed of a number of ASEs, one of which is always the *Association Control Service Element* (ACSE). This ASE has sole use of Presentation Service facilities for connection establishment and release. The ACSE augments the addressing mechanisms provided in the Presentation Layer by enabling an application-entity to be identified by a *application-entity-title* (AE-title) instead of a presentation-address.

An AE-title is a globally unambiguous name of the same general hierarchical form as X.400 mail addresses which were illustrated in [8]. The advantage of AE-titles over presentation-addresses is equivalent to the identification of people by names rather than telephone numbers. In due course the *OSI Directory Service* explained in [9] will be available to translate AE-titles into presentation-addresses. In the meantime it will be necessary to provide local mechanisms for this purpose. Registration procedures for AE-titles are currently being established by ISO in cooperation with all ISO member bodies. In the USA national registration procedures are currently being established by ANSI.

Abstract and Transfer syntaxes

Within the OSI mapping of the X Window System, the ACSE services are used to establish the connection between an X Window System client and server. Once the connection is established, the first request transmitted by the client is OPEN DISPLAY. This, and all subsequent messages between client and server, are carried by the P-DATA service. Normal termination by a CLOSE DISPLAY request and abnormal termination as a result of a KILL CLIENT request make use of the corresponding termination facilities of ACSE.

In contrast to the Transport Service which passes an uninterpreted byte stream transparently between its users, the Presentation Service passes messages known as *Presentation Data Values* (PDVs) which have a specific syntactic structure. This structure is specified by an abstract syntax that is chosen by the presentation-service-users through the exchange of ASN.1 Object Identifier values during connection establishment. An example of a PDV in the context of the X Window System is a BELL request with an 8-bit signed integer as a parameter; the parameter is used to specify the bell volume. The PDV in this form is abstracted from the coding used to indicate a BELL request and from the internal representation of signed integer values. Each end-system will have its own internal concrete syntax for the representation of PDVs. It is a function of the Presentation Layer to translate this internal representation to and from the transfer syntax used for transmission. The choice of transfer syntax is similarly negotiated during connection establishment, but it is negotiated privately by the two presentation protocol machines in the light of the abstract syntax that it is to represent.

The X Window SystemTM and ISO VT (*continued*)

The presentation-service-users are not party to this negotiation, neither are they aware of its result. Two phases of translation thus occur during the transmission process that are absent at the level of the Transport Layer.

Part I of the Standard provides both the abstract syntax and the semantics of the functions of the X Window System. Part II provides a transfer syntax. To enable these to be used in the OSI mapping, ANSI have registered ASN.1 Object Identifiers for both the abstract and transfer syntaxes. The Standard requires this particular transfer syntax always to be used for values specified in the X Window System abstract syntax.

Comparison of X and VT

When used with its OSI mapping, the X Window System provides a graphical OSI terminal standard just as ISO VT provides a character-oriented terminal standard. In both cases the human user interface is outside the scope of the respective standards. Both are capable of running on a wide range of platforms. They may thus be seen as competitors.

When the differences in capability and in required resources are examined, a different view emerges. The graphical terminal of the X Window System is clearly far more powerful than the character-oriented terminal of ISO VT. But a heavy price is paid in resource requirements. Every action by the human user at the X server has to be communicated to the X client for interpretation. Even a simple pointer movement through use of a mouse will generate a great deal of network traffic. In contrast ISO VT has facilities for the remote application to download local processing capability to the terminal. Quite complex editing operations may then be performed without any network traffic at all being generated until the final result is transmitted back to the application.

For the X Window System, reaction times acceptable to the human user require a high bandwidth. A Local Area Network (LAN) presents no problems in this respect, but operation over a Wide Area Network (WAN) and especially over a concatenation of such networks is likely to give rise to unacceptable delays. In contrast, ISO VT is designed with the problems of human user response particularly in mind. With downloaded local processing capability, the perceived response time is isolated from network transit times and concatenations of wide area networks cause no problems.

Cooperation between X and VT

Competition is not actually necessary. The power of the X Window System for local working can be combined with the rapid perceived response and local processing capability of ISO VT for remote working. The best of both worlds can thus be achieved, by providing the capability for ISO VT to run in a window on an X server. A gateway function to achieve this will consist of an Application Layer relay between an X client and a Terminal VT-user.

No additional functionality will be required in the terminal, which will be a standard X server. The gateway will be an application process that is invoked simply by opening a window on its X client. Communication between the X client and X server will be at high bandwidth. The gateway application process will then provide the facilities for establishment of a VT-association. This association may be to an application on a remote host that is connected to the gateway only through a concatenation of slow WANs.

Once the VT association is established, the gateway will be transparent to the human user. The remote application will be running in an X window that can be moved, resized and generally handled by the mechanisms of X. The local functionality of ISO VT will be provided within the gateway, and so will be subject only to the negligible transit delays of the local X connection. The X server may simultaneously have windows open to local X clients and to remote VT applications. Cooperation rather than competition should enable both standards to work together for mutual enhancement.

References

- [1] David Chappell, "Components of OSI: A taxonomy of the players," *ConneXions*, Volume 3, No. 12, December 1989.
- [2] Robert W. Scheifler, James Gettys & Ron Newman, "X Window System—C Library and Protocol Reference," Digital Press, 1988.
- [3] Graham Dixon, "Components of OSI: The Virtual Terminal ASE," *ConneXions*, Volume 5, No. 2, February 1991.
- [4] Bill Jolitz, "X Windows: More than Just a Pretty Face," *ConneXions*, Volume 4, No. 5, May 1990.
- [5] International Organization for Standardization, "Information processing systems: Open Systems Interconnection, Basic Reference Model," International Standard ISO 7498, 1984.
- [6] Kim Banker, "Components of OSI: The Session Service," *ConneXions*, Volume 3, No. 9, September 1989.
- [7] David Chappell, "Components of OSI: The Presentation Layer," *ConneXions*, Volume 3, No. 11, November 1989.
- [8] Julian Onions, "Components of OSI: The X.400 Message Handling System," *ConneXions*, Volume 3, No. 5, May 1989.
- [9] Steve Benford, "Components of OSI: The OSI Directory Service," *ConneXions*, Volume 3, No. 6, June 1989.
- [10] Wayne Dyksen & John T. Korb, "A Programmer's Overview of X," *ConneXions*, Volume 4, No. 10, October 1990.

GRAHAM DIXON is Director of Studies in Mathematics at Churchill College, Cambridge, England. He also acts on behalf of the College as a consultant on OSI. He is a member of the NIST OIW Special Interest Group on VT and of the EWOS Expert Group on VT. He received his Ph.D. degree in mathematics in 1965 and is the author of a book on Special Relativity published by Cambridge University Press.

Mark your calendar!

INTEROP 91 is only 5 months away. In a few weeks you should be receiving the INTEROP Advance Program with information about the INTEROP 91 conference sessions, tutorials, Birds Of a Feather sessions and more. Plan to join your fellow networking professionals in San Jose, October 7-11. For further information or to register, call us at 1-800-INTEROP or 415-941-3399, or fax us at 415-949-1779.



Resource Discovery and Related Research at the University of Colorado

by Michael F. Schwartz, University of Colorado—Boulder

Introduction

For the past three years, the Networked Resource Discovery Project at the University of Colorado, Boulder has explored a number of experimental means by which users can discover the existence of resources in a large internet environment. Example resources include network services, documents, retail products, current events, and people. This *resource discovery* problem is important because it is an enabling (and currently limiting) technology for the larger issue of supporting distributed collaboration, or the accomplishment of tasks through sharing resources among many interrelated individuals across administrative boundaries. Without the ability to discover resources of interest, users perceive only a very limited fraction of the full potential for sharing resources and collaborating with colleagues.

Goals

We impose three key goals on our approaches to resource discovery. First, we consider very large environments, spanning national or international networks. Such environments place stringent scalability requirements on the algorithms that can be used. Second, we want to avoid imposing artificial constraints on the resource space organization. Traditional directory services (such as the CCITT X.500 standard [1]) rely on hierarchical organization to achieve good scalability. Unfortunately, the organization of a hierarchy becomes convoluted as an increasingly wide variety of resources is registered, and requires users to understand how the (increasingly deeply) nested components are arranged. Finally, we wish to minimize the need for global administrative agreement over protocols, information formats, and organizational structures. While standards are helpful, it is difficult to specify standards that are both globally adopted and technologically current. Moreover, standards based on a hierarchical organization require a high degree of agreement over the organization of at least the upper levels of the tree. As an increasingly diverse collection of institutions contribute to the global information infrastructure, smooth evolution will require the ability to support multiple organizational structures, and to interoperate with a heterogeneous set of protocols and information formats.

During the course of our research, we have found that the techniques we have developed apply to a number of network problems beyond resource discovery and distributed collaboration, including network management, network integration, and network measurement. We overview a range of our research efforts below.

Netfind

The most mature and applied subproject we have completed to date is an Internet "white pages" directory tool called *netfind* [25]. Given the name of a user and a rough description of where the user works (e.g., the company name or city), *netfind* attempts to locate telephone and electronic mailbox information about that user. This research focuses on the ability to use a number of existing protocols and highly decentralized sources of relatively unstructured information. Using decentralized information avoids difficult problems of consistency and transfer of authority that are inherent in mechanisms that rely on building auxiliary databases to hold resource information, a point we will return to shortly. The ability to use simply structured information is important in heterogeneous, administratively decentralized environments, where global agreement about highly structured information formats is difficult to achieve.

The technique we have developed to support these characteristics involves building an understanding of the semantics of a particular resource discovery application into the algorithms that support searches. In the current case, the particular application is Internet white pages, and the semantically cognizant mechanism is as follows.

Approach

We begin with a database of “seed” data, which provides hints of potential machines to probe when a search is requested. This database is built by gathering information from the headers of USENET [17] news messages over time. These headers typically list the user name, organization name, city, and electronic mailbox for users who post messages. When a search is requested, the seed database is consulted to locate the names of a number of machines associated with institution keywords specified in the search request. Requests use the format “UserString InstString [InstString...],” where UserString identifies the user (typically by last name), and the conjunction of one or more InstStrings identify the institution where the user works. For example, a search could be requested for “schwartz university colorado” or “schwartz boulder.”

If the machines found in the seed database fall within more than three naming domains (an example of one domain being colorado.edu), the user is asked to select at most three domains to search. The *Domain Name System* (DNS) [11] is then contacted, to locate authoritative name server hosts for each of these domains. The idea is that these hosts are often central administrative machines, with accounts and/or mail forwarding information for many users at a site. Each of these machines is then queried using the *Simple Mail Transfer Protocol* (SMTP) [15], in an attempt to find mail forwarding information about the specified user. If such information is found, the located machines are then probed using the *finger* protocol [8]. The results from *finger* searches can sometimes yield other machines to search as well. Ten lightweight threads are used to allow sets of DNS/SMTP/*finger* lookup sequences to proceed in parallel, to increase resilience to host and network failures.

Because many different institutional keywords will lead to the same seed database records and Domain information, it is usually quite easy to “guess” keywords that will succeed for any particular search. Moreover, *netfind* can often find a user even if the remote site does not support all of the above protocols, or if some steps in the protocol sequence fail. For example, if *finger* is disabled because of security concerns, mail forwarding information may sometimes still be found. Or, if no mail forwarding information is found, *netfind* attempts to *finger* some of the machines matched from the seed database. Similarly, *netfind* can proceed without information about authoritative name servers. This ability to function in the presence of failures or partial remote protocol support is an example of a technique for supporting fault tolerant resource discovery without global agreement. We utilize this technique to a more significant extent in our network visualization project, described later in this article.

Netfind’s tolerance of partial remote protocol support allows it to locate information about a large proportion of Internet users. Measurements indicate that the scope of the directory is upwards of 1,147,000 users in 1,929 sites. This scope is significantly larger than other existing Internet directory services, which require that users register with an administratively centralized service (as with the Network Information Center *WHOIS* service [9]), or that special directory servers be run at many sites around the Internet (as with X.500).

continued on next page

Resource Discovery and Related Research (*continued*)

Netfind's ability to use highly decentralized information sources allows it to locate very timely information about users. Unlike services that use an auxiliary database that must be updated by a separate administrative procedure, *netfind* probes the machines on which users do their daily computing. To the best of the author's knowledge, all other white pages services (including WHOIS and X.500) depend on some form of auxiliary database. Populating and keeping such a database up-to-date are difficult tasks.

HNS

Using information where it naturally resides is a principle carried forward from our earlier *Heterogeneous Name Service* (HNS) work [18]. However, *netfind* uses much more decentralized information than the HNS did. The HNS was essentially a framework for supporting users in specifying the semantic operations needed to incorporate new auxiliary database-style name services into a global name service. Moreover, the HNS was used for mapping named objects to data about those objects (such as the network address of a host), rather than for discovering resources. More recently, Droms used an architecture similar to the HNS in his *Knowbot Information Service*, to provide a white pages service [5].

Netfind is in active use by researchers at approximately 50 institutions worldwide, and is being developed further commercially. Measurements of these users indicate that an average search uses approximately 137 packets. While this is more costly than a registration-style directory like X.500, we believe it is quite a reasonable price to pay for providing timely information without global cooperation, particularly when one considers the capacity of next-generation high speed networks.

Electronic Mail study

Another completed but less "tool oriented" project is a measurement study of global electronic mail communication patterns [23]. In this study, we sought to understand how people take part in distributed collaboration, and how distributed collaboration might be better supported. While the main form of distributed collaboration is currently electronic mail, the possibility exists for a much more significant degree of sharing. Electronic mail is primarily used in a point-to-point fashion, supporting the interchange of messages between pairs of individuals. We are interested in a more concurrent, symmetrical style of collaboration involving many participants. For example, a more powerful sharing mechanism could be modeled on the types of interactions that take place at conferences and other meetings, where people discuss issues collectively with people they had not previously met, but who are known to have closely related interests. Electronic "bulletin board" services such as USENET support a crude form of this collaboration, but their restrictive organizational structure (a small number of relatively statically defined interest groups) and means of information distribution (unsequenced, full scale broadcast) do not readily facilitate high quality collaboration.

Social networks

As a point of departure, we became interested in the organization of human social networks. Such networks use a non-hierarchical organizational structure that scales well. Rather than forming contacts with each other based on a hierarchy, people often establish more direct "networks," by contacting knowledgeable intermediaries who can quickly refer them to other relevant people, cutting across bureaucratic boundaries. For example, by contacting a computer science professor or network manager, someone interested in high-speed networking technology can quickly meet other people who share this interest.

These people can, in turn, introduce the person to others who perhaps more closely share his/her particular interests. At the same time, the newcomer can be instrumental in pointing out individuals who share other interests with the people he/she meets.

To study these organizational properties, we collected mail logs from 15 sites around the U.S. and Western Europe for two months. This data constituted a graph containing approximately 50,000 users in 3,700 different sites around the world. Applying graph theoretic analyses yielded a number of insights about how users collaborate by electronic mail. We found that the average path between people is short, which is a statistically rendered version of the small diameter postulated by the so-called “small world” phenomenon [27]. This property indicates that the graph can support rapid information dissemination. Furthermore, we found that the graph edges are highly redundant, indicating that the graph can support reliable information dissemination. These properties are highly sought in computer networks, yet arise naturally in human social networks.

SSGs

From the perspective of distributed collaboration, an even more interesting characteristic of human social networks is their flexible organizational structure. Rather than forcing all internode relationships to conform to a hierarchy, the graph structure allows individuals to be related to one another through multiple groupings that we call *Specialization Subgraphs* (SSGs). An SSG is a subset of nodes that share common attributes, and that has a small diameter. As an example, in a graph of relationships among people, one SSG could connect individuals based on a shared interest in a particular computer science speciality, a second SSG could connect individuals based on shared responsibilities at a place of employment, and a third SSG could group individuals based on shared cultural/recreational interests. Any individual can belong to many different SSGs, and can search for information about a particular topic by consulting the appropriate SSG.

To measure properties of SSGs, we developed an algorithm to cluster individuals by shared interests without access to the contents of the mail messages, by computing properties of the mail interconnection graph using traffic analysis techniques [2]. We found it necessary to apply the traffic analysis to a subgraph derived by a graph reduction technique that eliminated “noise” caused by our statistical sampling. The combination of these two techniques allowed us to derive lists of individuals who share interests with one or more specified individuals. (We ran the computation specifying a number of individuals whose interests were known to us to validate the procedure.) This clustering algorithm has potential applications for distributed collaboration (discussed further in the Section “Internet Resource Mapping/Discovery Project” below), as well as privacy implications for electronic mail. Moreover, applying this algorithm to each of a randomly chosen subset of 500 nodes within the graph provided measurements of how people collaborate, indicating the existence of a large number of different but heavily interrelated groupings of individuals based on shared interest, and underscoring the importance of supporting a number of different organizational structures for distributed collaboration.

Probabilistic Yellow Pages protocols

Another project on which we have made substantial progress is a study of probabilistic algorithms that construct and search a resource graph that supports attribute-based (“yellow pages”) specifications [20, 22].

continued on next page

Resource Discovery and Related Research (*continued*)

For this project we assume it is desirable to find a small number of instances of a moderately large class of objects. For example, in searching for a supplier of a particular piece of computer hardware, finding 5 out of 100 suppliers in a metropolitan area would often suffice. We also assume that it is acceptable to return different answers to the same query across search sessions. If consistent responses to queries are desired, one could build a front-end user interface that cached results, and provided identical responses across search sessions.

Brokers

Based on these assumptions, we designed and simulated a protocol to support a set of agents in organizing and searching the resource space. Agents maintain pointers to sources of resource information, and access these sources via intermediary *brokers*, which enforce the access control policies and encapsulate the heterogeneity of the information repositories. While agents are intended to be part of the network infrastructure, each broker belongs to the organization whose resource information it exports.

While brokers are an important part of the model, the main focus of our research is on the agent protocols for organizing and searching the resource space. Rather than having an administrative body specify how the space is organized, agents organize the space dynamically, according to what resources exist and the types of searches that users make. Agents use a probabilistic *Sparse Diffusion Multicast* primitive to disseminate information about resources at uniformly distributed, randomly chosen nodes around the network, and likewise to route search requests randomly around the network. Sparse Diffusion Multicast can be implemented with simple modifications to Cheriton and Deering's Internet Multicast protocol [4].

Randomly disseminating resource information is intended to place the information within a reasonably small neighborhood of any agent in the network, so that during searches it is likely that the information can be found using simple random probes. Since the types of resources that exist and the searches users request are not random, a cache management policy is used to prefer graph edges between agents that maintain related information, to form specialization subgraphs. Using this policy, a search initiated at a random agent may cause some random search behavior at the start of the search, until a member of an appropriate specialization subgraph is reached. If such a subgraph is reached, searches proceed in a more directed fashion. If a user continues to use the same agent, over time that agent will maintain pointers to sources of the type of information for which that user often searches.

Our results to date indicate that this approach can support a non-hierarchical resource space for an environment roughly the size of a country, with several thousand sites participating in resource registration and searches. We are currently extending these results to model various aspects of the protocol in more detail.

The probabilistic nature of this approach also supports fair access among competing information providers, an important issue in commercial environments such as the U.S. telecommunications industry [7] or computer reservation systems [6].

Internet Resource Mapping/Discovery project

A project currently under way is an effort to support resource discovery in decentralized environments of such scale that the resource space cannot be completely organized.

In such an environment, mechanisms are needed to support incremental organization of the resources, based on the efforts of many geographically distributed individuals, and a range of different information sources of varying degrees of quality. Our approach to this problem is to use mechanisms that “tap into” existing network protocols and information sources to provide an immediately useful tool (much as netfind did), supplemented by mechanisms that allow users to superimpose additional organization on the resource space in an incremental fashion. As a concrete test case, we are developing a prototype that focuses on public Internet archive sites accessible via the “anonymous” FTP [16]. This is an interesting test case, because it encompasses thousands of administratively decentralized sites containing a collection of resources of considerable practical value.

The part of this effort associated with tapping into existing infrastructure is now basically complete [21]. In this prototype, three levels of information quality are supported. At the highest level, resources are described using archive-site-resident databases, with individual resources described according to their conceptual roles. Below that, per-user and per-user-site caches are maintained, to record resources that have been found by individual users during their explorations. At the lowest level, the system scans USENET electronic bulletin board articles using a simple set of heuristics to recognize announcements about public archive sites, to provide a simple keyword-based index of resources throughout the Internet.

Views

To support users in superimposing additional organization on the resource space, we are currently developing a system architecture that allows any individual or group of users who share common interests to build a structure (called a *view*) that superimposes organization on the resource space according to their particular interests. For example, a group interested in graphics might build a view that organizes the world according to PostScript, Tools, Window Systems, Images, and Discussions, with pointers to network accessible resources of these various types nested into this structure. A view is intended to be a simple structuring mechanism for loosely integrating an administratively decentralized pool of resources, with properties much like those of specialization subgraphs. Views can include pointers to parts of other views, so that related interest groups (such as people interested in operating systems and people interested in data communications) can cross reference each others' views. Views are not constrained to be hierarchical, although in practice many links will probably be tree-like.

Because an arbitrary number of different views may coexist, our research investigates how to support users in discovering, constructing, and sharing views. We consider several problems. Views must be replicated and distributed, to improve availability and performance. Views must be secured against accidental or malicious modification. It must be easy to update and reorganize views. Finally, it must be easy to search for resources.

Searches

We place particular focus on supporting searches. Our approach is to build a simple flat index of each view. While flat searching has not worked well in some situations involving large scale (such as library information systems), searching a view in this manner should work well because any particular view will be fairly small and highly focused in scope. Because of this, it should be relatively easy to “guess” appropriate keywords for searching a view, without the difficulty of keywords matching many unrelated subjects.

continued on next page

Resource Discovery and Related Research (*continued*)

Moreover, the underlying space in a view is structured. Hence, a user could examine a subtree in a view once a match occurs, unlike the corresponding situation with the flat underlying spaces supported by information retrieval systems.

Of course, the problem remains that a user must discover appropriate views to search when trying to find a resource. To address this problem, an area of future research we intend to pursue involves experimenting with a means of automatically interrelating views, based on the interest clustering algorithm developed in connection with our electronic mail study. Doing so will provide a dynamically evolving set of links between related views, and allow users to search for resources without having to know what views exist.

We will deploy the prototype at a number of Internet sites, and enlist the participation of users in building a distributed map of the resources available via anonymous FTP. Sites interested in participating in this effort when the software becomes stable are encouraged to contact the author.

Network Visualization Project

Another project currently under way involves using resource discovery techniques to support visualization of characteristics of the global Internet, such as topology, congestion, routing, and protocol usage [24]. As with netfind and the anonymous FTP prototype, we will use a number of protocols and information sources, to support discovery in the absence of global agreement on any one protocol or information source. For the visualization project, however, we will use a much more extensive collection of protocols and information sources, including the *Address Resolution Protocol* [13], the *Internet Control Message Protocol* [14], the *Domain Name System* [11], the *Simple Network Management Protocol* [3], and a dozen others. Essentially, this project is exploring the extent to which one can integrate a heterogeneous, administratively decentralized network by using resource discovery techniques.

We will deploy a collection of servers around the Internet, each of which periodically executes a set of discovery protocols to maintain information about its local internet. An X-Window user interface will allow users to connect to servers and browse the state of the Internet. Servers will cache information and support queries that allow partial retrieval of remote information, for example retrieving information that has changed since the last update to the local cache contents. Servers will also be parametrized so that system administrators may chose what discovery protocols will execute and their scheduling frequency, according to their perceptions of the relative importance of network loading, timeliness, and completeness of the data.

Wide Area Demand Resource Distribution Project

We have recently begun work on a mechanism for widely distributing files without broadcasting them to all nodes on the network (as is the case with the *Network News Transfer Protocol* [10]), and without causing files to be retrieved multiple times across individual network links (as is the case with anonymous FTP). The basic idea, suggested by Phil Karn of Bell Communications Research, is to distribute files in response to requests for them, caching them at intermediate nodes along a dynamically-developed spanning tree of the Internet. Doing so can potentially reduce the load on the Internet substantially, since FTP currently accounts for 45% of the bytes transmitted on the NSFNET backbone [12]. Moreover, this mechanism could underlie a network-transparent mode of resource sharing, whereby resources are named independently of the particular hosts that hold them.

This technique will be best supported by an Internet Multicast mechanism [4], but will also function without such specialized support. A number of issues must be addressed in this project, including accountability of cache contents (to ensure that intermediary sites do not accidentally or maliciously modify file replicas), cache consistency (to ensure that the most recent version of a distributed file is retrieved), and the proper level of file fragmentation, in response to alternative routing of file contents across the dynamically routed Internet.

Summary

This article presents a brief overview of a number of projects exploring techniques to support resource discovery and distributed collaboration. We focus on techniques appropriate for national or international networks, and the concomitant issues of scalability, administrative decentralization, and supporting natural organization. Our efforts involve wide area distributed prototypes and measurement studies, using the global TCP/IP Internet as an experimental testbed. In addition to resource discovery and distributed collaboration, the techniques we have developed also apply to a number of related problems, including network integration, network management, and, more generally, to supporting globally distributed applications [26].

Supporting resource discovery raises some difficult issues concerning privacy of information. While security mechanisms may be imposed to preserve privacy in some cases, in many cases such mechanisms are either difficult to provide, or of questionable merit. We believe that privacy is essentially a social issue, and as such requires careful consideration about the policies that will manage the technical solutions, rather than an emphasis simply on technical solutions to security problems. Moreover, we believe that the best way to understand the tension between privacy and resource discovery is to explore the issues raised by building and deploying resource discovery prototypes.

Interested readers may obtain copies of many of the project related papers by anonymous FTP from `latour.colorado.edu`, in the directory `pub/RD.Papers`, or by contacting the author.

References

- [1] CCITT, "The Directory, Part 1: Overview of Concepts, Models and Services," ISO DIS 9594-1, CCITT X.500, 1988.
- [2] L. D. Callimahos, "Traffic Analysis and the Zendian Problem," Aegean Park Press, 1989.
- [3] J. Case, M. Fedor, M. Schoffstall & C. Davin, "A Simple Network Management Protocol (SNMP)," RFC 1098.
- [4] D. R. Cheriton & S. E. Deering, "Multicast Routing in Datagram Internetworks and Extended LANs," *ACM Transactions on Computer Systems*, 8(2), May 1990.
- [5] R. E. Droms, "Access to Heterogeneous Directory Services," Proc. of the *InfoCom* conference, June 1990.
- [6] D. Gifford & A. Spector, "The TWA Reservation System," *Communications of the ACM*, 27(7), July 1984.
- [7] H. H. Greene, United States of America, Plaintiff, v. Western Electric Company, Inc., et al., Defendants. Civil Action No. 82-0192, U.S. District Court, District of Columbia, March 1988, "Triennial review of Modified Final Judgement divesting AT&T of the Regional Bell Operating Companies."
- [8] K. Harrenstien, "Name/Finger," RFC 742.
- [9] K. Harrenstien, M. Stahl & E. Feinler, "NICName/Whois," RFC 954.

continued on next page

Resource Discovery and Related Research (*continued*)

- [10] B. Kantor & P. Lapsley, "Network News Transfer Protocol—A Proposed Standard for the Stream-Based Transmission of News," RFC 977.
- [11] P. Mockapetris, "Domain Names—Concepts and Facilities," RFC 1034.
- [12] NSF Network Service Center, "Internet Resources Guide," Available via anonymous FTP from `nnsd@nnsd.nsf.net`.
- [13] D. C. Plummer, "An Ethernet Address Resolution Protocol—Or—Converting Network Protocol Addresses to 48-bit Ethernet Address for Transmission on Ethernet Hardware," RFC 826.
- [14] J. Postel, "Internet Control Message Protocol," RFC 792.
- [15] J. Postel, "Simple Mail Transfer Protocol," RFC 821.
- [16] J. Postel & J. Reynolds, "File Transfer Protocol (FTP)," RFC 959.
- [17] J. S. Quarterman & J. C. Hoskins, "Notable Computer Networks," *Communications of the ACM*, 23(10), October 1986.
- [18] M. F. Schwartz, J. Zahorjan & D. Notkin, "A Name Service for Evolving, Heterogeneous Systems," *Operating Systems Review* 21(5).
- [19] M. F. Schwartz, "Autonomy vs. Interdependence in the Networked Resource Discovery Project," ACM SIGOPS European Workshop, September 1988.
- [20] M. F. Schwartz, "The Networked Resource Discovery Project," Proc. of the IFIP XI World Congress, August 1989.
- [21] M. F. Schwartz, D. R. Hardy, W. K. Heinzman & G. Hirschowitz, "Supporting Resource Discovery Among Public Internet Archives Using a Spectrum of Information Quality," Technical Report CU-CS-487-90, University of Colorado, September 1990.
- [22] M. F. Schwartz, "A Scalable, Non-Hierarchical Resource Discovery Mechanism Based on Probabilistic Protocols," Technical Report CU-CS-474-90, University of Colorado, June 1990.
- [23] M. F. Schwartz & D. C. M. Wood, "A Measurement Study of Organizational Properties in the Global Electronic Mail Community," Technical Report CU-CS-482-90, University of Colorado, August 1990.
- [24] M. F. Schwartz, D. H. Goldstein, R. K. Neves & D. C. M. Wood, "An Architecture for Discovering and Visualizing Characteristics of the Global Internet."
- [25] M. F. Schwartz & P. G. Tsirigotis, "Experience with a Semantically Cognizant Internet White Pages Directory Tool," *Journal of Internetworking Research and Experience*, 2(1), March 1991.
- [26] M. F. Schwartz & P. G. Tsirigotis, "Principles for Supporting Globally Distributed Applications."
- [27] J. Travers and S. Milgram, "An Experimental Study of the Small World Problem," *Sociometry*, 32(4), pp. 425–443, 1969.

MICHAEL SCHWARTZ received his B.S. (1982) from UCLA, and his M.S. (1985) and Ph.D. (1987) from the University of Washington. He is currently an Assistant Professor of Computer Science at the University of Colorado. His research focuses on issues raised by national and international networks and administratively decentralized systems, with particular emphasis on resource discovery and distributed collaboration. He can be reached as `schwartz@latour.colorado.edu`.

Book Reviews

Ed.: This month, we bring you reviews on two recently published books on computer and network security:

- Security Study Committee (David D. Clark, Chairman), National Research Council, *Computers at Risk: Safe Computing in the Information Age*, National Academy Press (Washington, D.C., 1991). ISBN 0-309-04388-3, 320 pages, paperback. Available in bookstores or by calling 1-800-624-6242.
- Peter J. Denning, Ed., *Computers Under Attack: Intruders, Worms, and Viruses*, Addison Wesley (Reading Mass., 1990). ISBN 0-201-53067-8, 554 pages, paperback.

Computers at Risk

"We are at risk..." So begins a report of the National Research Council, a research arm of the National Academy of Sciences, on the subject of computer security. The report is the result of a year-long effort chaired by the ubiquitous David D. Clark of MIT.

The committee includes such notables as Butler Lampson, architect of DEC's security framework, and Peter Neumann of SRI International, noted for his in-depth list of security problems. The committee also included Stephen Kent of BBN, author of the RFCs for the Internet *Privacy Enhanced Mail* (PEM) prototype.

Well-written

In contrast with most treatments of the subject of computer security, this report is extremely well-written. Not only is the report readable, it is significant. The Clark committee pursued a broad mandate from DARPA, the originator of the study, to look at a "national research, engineering and policy agenda to help the United States achieve a more trustworthy computing technology base by the end of the century."

Quite a charter. With a mandate like that you have to go beyond just recommending longer passwords and the committee does in fact come up with an extremely comprehensive policy.

Security is more than just protecting your own assets. It is a fundamental aspect of being a good corporate citizen. Banks protect against breakins to protect their own assets, but also because they have an obligation to protect the sensitive information they keep for their clients.

Holistic concern

Likewise, members of the Internet should protect their own system out of self-interest, but also out of a broader mandate to protect their neighbors. Security is a broad, holistic concern requiring the participation of all. In this sense, many of our corporate citizens have failed miserably. Take the credit reporting bureaus that allow hackers repeated access to their systems because increased protection is not directly profitable to themselves. Or, take the case of the wide-open Internet systems that allow hackers to penetrate other systems.

The root cause of these problems is the concept of an externality: a cost not borne by the person taking the action and thus borne by others. The Clark committee realizes this and proposes a broad national agenda to strengthen security systems.

GSSP

The basic recommendation in this book is the promulgation of *Generally Accepted System Security Principles*, (GSSP) a broad-based set of principles that everybody can accept, akin to the *Generally Accepted Accounting Principles* (GAAP) in the accounting profession.

continued on next page

Book Reviews (*continued*)

Information Security Foundation

Accompanying the GSSP, the committee recommends many other steps. Some are to be expected from a committee of researchers: fund more research, for example. Other conclusions are ones that the industry has long clamored for such as lifting export restrictions on DES.

Perhaps the most controversial recommendation of the committee is to start a new institution, an *Information Security Foundation*. This new non-profit, non-governmental group would help develop the GSSP, educate the public, and perform a variety of other important tasks.

What makes this report so important is not the final recommendation for yet another coordinating body, but the analysis that leads up to it. An in-depth discussion of what security is makes the report must reading for anybody who has responsibility for computer systems.

The committee goes through technical means of providing security, but quickly makes the point that the most secure system is useless if run by poorly-educated people. Security is more than just hardware: it is procedures, education, and most importantly, a continuing focus by senior management.

A large part of the report is devoted to the development of software. Proper design, verification, and testing allow software to be designed secure instead of shut after the fact. Particularly important here is the use of high-level software: programs written in high-level languages are more secure than those hacked together in assembly language with unforeseen holes.

Granted, applications in high-level software only make sense if the underlying environment is a trusted computing base. If you make a high-level system call, it would be nice to know that an attacker is unable to tunnel into the system at lower levels and subvert the call.

Computers Under Attack

Computers Under Attack, edited by Peter Denning, is a very different book—the detective movie-like cover of a man in a trenchcoat gives one the first clue that this book is meant to tantalize the reader instead of setting a national policy.

The book is a compilation of 40 articles by the editor-in-chief of the venerable *Communications of the ACM*. The bulk of it is devoted to a series of articles that explain breakins, worms, and viruses.

All the classics

This book has all the classics. Well-known incidents like the IBM Christmas Card are detailed: when you read the card, it would automatically send itself to everybody on your mailing lists, quickly swamping the network it attacked.

The first part of the book may be useful to some, but will have little interest to most readers of *ConneXions*. It contains classic articles explaining what an internet is, including the classic “Notable Computer Networks” by Quarterman and Hoskins. The reader wanting more of this information is directed to Quarterman’s recent book, *The Matrix: Computer Networks and Conferencing Systems Worldwide* (Digital Press, ISBN 1-55558-033-5, 1990).

In section Two through Four, the book takes off. It starts with intruders, including Brian Reid’s discussion of how he was unable to track down a breakin and (of course) the Clifford Stoll attempt to recover 75 cents by calling the FBI. (You can get the full story by reading Stoll’s own account in his book *The Cuckoo’s Egg*).

The book then goes on with a large set of articles detailing the Morris *Internet Worm*. Particularly interesting is Eugene Spafford's detailed dissection of how the worm operated. The virus section includes, among more general articles, a detailed look at the Pakistani virus.

Issues

The last two sections of this book are a potpourri of articles on legal, social, and ethical issues. Particularly interesting is the interview by Dorothy Denning, a DEC researcher by an anonymous Cyberpunk using the handle "Frank Drake." After Dorothy Denning answered the questions posed to her (by e-mail of course) she sent a TURN message and started asking her own questions. The resulting dialogue is fascinating.

Peter Denning says in his forward that this book should be considered like a room full of people. The book contains no new material of any significance, but just putting all those people in one room is useful.

Recommended

If you need to develop a security policy, start with the Clark report. If you want to scare your boss into accepting that policy, read *Computers Under Attack*. Both are recommended. —Carl Malamud

CARL MALAMUD is the author of *DEC Networks and Architectures* (McGraw-Hill, 1989), *INGRES* (Van Nostrand Reinhold, 1989), *Analyzing Novell Networks* (Van Nostrand Reinhold, 1990), and *Analyzing DECnet/OSI Phase V* (Van Nostrand Reinhold, 1991). He is currently writing *STACKS*, the official book of INTEROP 91. The official candy bar of INTEROP 91 has not yet been chosen. Carl can be reached as carl@malamud.com.

The 5th Wave

By Rich Tennant

The computer virus crept silently from network to network, until it found its way into the cafeteria vending machines.



CONNEXIONS

480 San Antonio Road
Suite 100
Mountain View, CA 94040
415-941-3399
FAX: 415-949-1779

FIRST CLASS MAIL
U.S. POSTAGE
PAID
SAN JOSE, CA
PERMIT NO. 1

ADDRESS CORRECTION
REQUESTED

CONNEXIONS

EDITOR and PUBLISHER Ole J. Jacobsen

EDITORIAL ADVISORY BOARD

Dr. Vinton G. Cerf, Vice President,
Corporation for National Research Initiatives

A. Lyman Chapin, Chief Network Architect,
BBN Communications

Dr. David D. Clark, Senior Research Scientist,
Massachusetts Institute of Technology

Dr. David L. Mills, Professor,
University of Delaware

Dr. Jonathan B. Postel, Communications Division Director,
University of Southern California, Information Sciences Institute

CONNEXIONS

Subscribe to CONNEXIONS

U.S./Canada ☐ \$150. for 12 issues/year ☐ \$270. for 24 issues/two years ☐ \$360. for 36 issues/three years

International \$ 50. additional per year (Please apply to all of the above.)

Name _____ Title _____

Company _____

Address _____

City _____ State _____ Zip _____

Country _____ Telephone () _____

☐ Check enclosed (in U.S. dollars made payable to CONNEXIONS).

☐ Visa ☐ MasterCard ☐ American Express ☐ Diners Club Card # _____ Exp. Date _____

Signature _____

Please return this application with payment to: CONNEXIONS

Back issues available upon request \$15./each
Volume discounts available upon request

480 San Antonio Road, Suite 100
Mountain View, CA 94040 U.S.A.
415-941-3399 FAX: 415-949-1779